Four-Dimensional Processing of Deformable Cardiac PET Data

Gregory J Klein
Center for Functional Imaging, Lawrence Berkeley National Laboratory
University of California, 1 Cyclotron Road, Berkeley, CA 94720

Abstract

A 4D deformable motion algorithm is described for use in the motion compensation of gated cardiac PET. The algorithm makes use of temporal continuity and a non-uniform elastic material model to provide improved estimates of heart motion between time frames.Temporal continuity is enforced by two methods. First, a prediction motion field is calculated using an assumption of constant particle velocity between time frames. Second, incremental motion fields are used as initializations for the estimation of motion fields between distant time frames. The prediction fields augment the algorithm's ability to estimate motion between noisy adjacent time frames, and the concatenated incremental motion fields improve estimation for large deformations. The estimated motion fields are used to establish a voxel correspondence between volumes, and produce a motion compensated composite volume.

1. Introduction

The left ventricle of the heart deforms by thickening and twisting as it beats during the cardiac cycle. For imaging devices which require longer than the time of the cardiac cycle to acquire data, like positron emission tomography (PET), this deformation induces blur proportional to the motion of the cardiac tissue. A method to combat the blurring is to perform what is called a gated cardiac acquisition. In this method, data are acquired over many cardiac cycles, and are divided into separate storage locations based on the time since the last cardiac R-wave. Typically, gated cardiac PET acquisitions divide each cardiac cycle into 50-100 msec time frames and then sum the data into the appropriate time frame over all cardiac cycles for an acquisition period ranging from 5 to 60 minutes. When each data set is reconstructed into an image volume, the images show the heart motion captured as a set of "freeze frames." Unfortunately, because the data have been divided into many different time frames, each reconstruction appears rather noisy due to a lack of statistics. We present here a summing technique which improves the contrast to noise characteristics of the images while compensating for the blur due to cardiac contractile motion. Each time frame is deformed using a non-rigid transformation so that the shape of the heart in the deformed volume matches a reference shape at end-diastole. By summing together the deformed volumes, the contrast to noise properties of the composite volume improve while the motion blur is reduced.

The approach presented here is unique from past efforts in modeling deformable motion in that it uses the temporal information present in the four-dimensional (4D) dataset directly to estimate a dense Lagrangian motion field. Recognizing that the cardiac contraction produces a cyclic, smooth motion between frames enables the algorithm to better estimate the motion than would an algorithm using only two time frames in isolation.

2. Related Work

A large body of work now exists that describes different techniques for deforming a three-dimensional (3D) voxel volume to match features in another 3D volume. Nearly all algorithms are composed of two main components. First, an image matching criterion serves as a driving force to push deformed image features into correspondence, and second, a regularization criterion prevents deformations which are physically unlikely or meaningless. For example, Song and Leahy [1] and later Zhou [2] proposed a 3D extension of Horn's [3] two-dimensional optical flow algorithm as a method to calculate the motion at each voxel, or a *motion field*, between ultrafast computed tomography (CT) sequences of the heart. Their method assumed that voxel values corresponding to the same physical features in two time frames conserve their value, much like the brightness constraint in Horn's 2D formulation. Therefore an image difference between two time frames could serve as an driving force. Regularization of the motion field was carried out by imposing a general smoothness and divergence-free constraint on the field.

Bajcsy and others were among the first to incorporate a more realistic material model as a regularization constraint [4, 5]. These authors used a linear elastic model to

constrain the deformation for brain datasets being warped to match a common atlas. Christensen built on this work, incorporating a viscous fluid model that could better model large deformations [6]. Others have combined material models implemented in a voxelized volume with surface-based methods [7, 8, 9]. Still others have found success using simpler regularization techniques such as a 9-parameter globally linear model [10], or low-dimensional global polynomial or spline-based deformations [11, 12, 13].

Techniques exploiting the 4D continuity of image sequences have been reported in the 2D optical flow literature for some time now [14,15,16]. However, perhaps because many deformable motion algorithms have been applied to the matching of different brain datasets, it is not surprising that these techniques did not consider 4D methods. Some have proposed 4D methods for tracking cardiac features, but these techniques have been used primarily in surface, not voxel-based deformation algorithms [17, 18]. Recently, encouraging results have also been obtained for tracking phase contrast cardiac magnetic resonance (MR) data using Kalman filtering [19] and Fourier spatiotemporal constraint models [20, 21]. 4D techniques have also been used to track features in cardiac tagged MR data [22] Though the 4D aspects of these techniques have primarily been applied to a relatively sparse set of data points, they may be applicable to the dense Lagrangian motion field description used here.

Due to its low temporal and spatial resolution compared to CT or MR imaging, cardiac PET and single photon emission computed tomography (SPECT) imagery pose a particularly difficult problem for the application of motion tracking. Though gated SPECT data has been used routinely for estimation of wall motion or ejection fraction, currently available motion analysis techniques typically focus on characterizing global motion parameters instead of estimating voxel-by-voxel point correspondences [23,24,25,18]. The present work will describe an algorithm for using 4D information to augment the estimation of a dense motion vector field for cardiac PET data.

3. METHOD

3. 1. General Cost Function

The approach described here is an extension of our previous 3D algorithm which uses a non-uniform regularization function inspired from elastic material models [26,27]. The goal of the algorithm is to find a dense Lagrangian vector field describing the motion of each voxel between a source volume and a reference volume. This vector field, which we call the motion field, is used together with the source volume and a forward voxel sampling scheme to create a deformation volume closely matching the reference. Because numerous motion fields could exist that would result in a deformation volume similar to the reference, but representing a physically implausible deformation, the motion field is regularized by an energy function constraining the source volume as if it was a physical volume composed of elastic materials being deformed by external forces. By choosing elastic parameters that describe the physical properties of the heart and adjacent tissue, estimation of physically implausible motion fields is prevented. Basically, then, the algorithm finds an appropriate balance between image matching forces and elastic constraining forces.

The algorithm is described as follows. Define a 4D voxel volume, $f(\boldsymbol{r}, t)$, where $r = (x, y, z)$ represents the voxel index and $t = \{1, 2, \ldots, N\}$ is the time index. A dense Lagrangian motion field describing the motion from time frame t_j to t_i is defined as $\boldsymbol{m}(\boldsymbol{r}, t_i, t_j) = (u(\boldsymbol{r}, t_i, t_j), v(\boldsymbol{r}, t_i, t_j), w(\boldsymbol{r}, t_i, t_j))$. The deformed volume of $f(\boldsymbol{r}, t)$ is defined as $\hat{f}(\boldsymbol{r}, t) = f(\boldsymbol{r} + \boldsymbol{m}, t)$. With these definitions, we can express an image matching error term, $e_I(\boldsymbol{r}, t_i, t_j)$, between a reference image at time, t_i and a source image at time t_j given by

$$e_I(\boldsymbol{r}, t_i, t_j) = \gamma_I[f(\boldsymbol{r}, t_i) - \hat{f}(\boldsymbol{r}, t_j)]^2 \quad (1)$$

where γ_I is a scalar weighting term.

A non-uniform elastic regularization term for the motion field based on a linear elastic strain energy model is given by

$$e_S(\boldsymbol{m}(\boldsymbol{r}, t_i, t_j)) = \frac{\lambda(\boldsymbol{r}, t_j)}{2}(u_x + v_y + w_z)^2 + \quad (2)$$
$$\mu(\boldsymbol{r}, t_j)(u_x^2 + v_y^2 + w_z^2) + \frac{\mu(\boldsymbol{r}, t_j)}{2}(u_y^2 + u_z^2 + v_x^2 + v_z^2 + w_x^2 + w_y^2 +$$
$$2u_y v_x + 2u_z w_x + 2v_z w_y)$$

where λ and μ are elasticity terms called the Lamé constants, and where derivatives of the motion field are denoted as $u_x = du/dx$. Though it is known that cardiac tissue exhibits nonlinear, anisotropic elastic properties that depend not only on tissue orientation, but also the muscle fiber direction and other factors [28, 29], we make a simplifying assumption that it may be modeled as a locally isotropic model described by just two terms, λ and μ. The drastic difference in elastic properties between the cardiac tissue and adjacent regions, such as the blood pool inside the left ventricle and lung tissue, is modeled by using a bivalued vector field for the Lamé constants. That is, voxels classified as cardiac tissue are modeled by one pair of λ and μ values which enforce the incompressible nature of this tissue. Voxels not belonging to the heart are characterized by Lamé constants which allow compression and which allow the material to be stretched with less force.

It will be assumed in this paper that an adequate segmentation of the left ventricle may always be obtained. Because cardiac PET data results in images with the left ventricle as the principal feature in the field of view, a rough segmentation of the heart suitable for this nonuniform elastic description may usually be obtained by smoothing each image volume considerably in the spatial domain and then using a simple thresholding operation. It is acknowledged that this simple segmentation operation may not be adequate for all cardiac datasets, especially for diseased hearts, where perfusion defects result in regionally reduced voxel intensities. For these cases, more powerful segmentation techniques are available [24,25].

A third energy term is derived from the smooth 4D nature of the deformation. A prediction motion field $\boldsymbol{m}_p(\boldsymbol{r}, t_i, t_j)$ computed from the temporally adjacent time frames is used to form the error term

$$e_p(\boldsymbol{m}(\boldsymbol{r}, t_i, t_j)) = \gamma_P[\boldsymbol{m}(\boldsymbol{r}, t_i, t_j) - \boldsymbol{m}_p(\boldsymbol{r}, t_i, t_j)]^2 \qquad (3)$$

The Lagrangian description of voxel motion allows us to view each point in the image volume as a particle moving at a specific velocity. The prediction field assumes that the velocity remains the same between time frames.

The overall estimation problem is to find a motion field consistent with elastic material properties that best matches the deformed volume to the reference volume via a minimization of:

$$E_{tot} = \sum_{\boldsymbol{r}, t} [e_I(\boldsymbol{r}, t_i, t_j) + e_S(\boldsymbol{r}, t_i, t_j) + e_P(\boldsymbol{r}, t_i, t_j)] . \qquad (4)$$

Solution to the problem is obtained via a Taylor expansion of $f(\boldsymbol{r} + \boldsymbol{m}, t_j)$ around $\boldsymbol{m}(\boldsymbol{r}, t_i, t_j)$ in a fashion similar to the technique originally proposed by Zhou [2]. During the minimization process, the true motion field, $\boldsymbol{m}(\boldsymbol{r}, t_i, t_j)$, is assumed to be composed of the current estimate, $\tilde{\boldsymbol{m}}(\boldsymbol{r}, t_i, t_j)$, plus a small delta motion field, $\delta\boldsymbol{m}(\boldsymbol{r}, t_i, t_j)$ that is to be computed. Expressing the overall cost function in terms of the Taylor expansion, $\hat{\boldsymbol{f}}(\boldsymbol{r}, t_j) = \boldsymbol{f}(\boldsymbol{r} + \tilde{\boldsymbol{m}}, t_j) - \nabla \boldsymbol{f}(\boldsymbol{r} + \tilde{\boldsymbol{m}}, t_j)\delta\boldsymbol{m}$ and substituting the expression, $\tilde{\boldsymbol{m}} - \delta\boldsymbol{m}$, for $\boldsymbol{m}$ results in a quadratic functional in $\delta\boldsymbol{m}(\boldsymbol{r}, t_i, t_j)$ that can be minimized via the calculus of variations [30]. The resulting Euler-Lagrange equations are solved using finite differencing techniques and a conjugate gradient method. For the simulations presented in this paper, Neumann boundary conditions were enforced [1].

3. 2. 4D Processing

Ultimately, it is desired that the image data in all time frames be deformed so that they match the shape of the heart at a single reference time. We choose end-diastole as the reference time. Here, the left ventricle is relaxed and the ventricular chamber is at its maximum size. Unfortunately, we have found in our prior efforts using only two

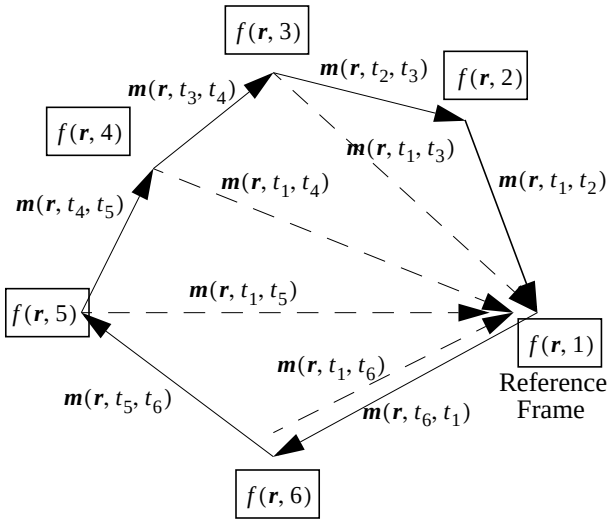

Fig. 1. 4D Motion Field Estimation. Motion estimation is begun by estimating adjacent time frames so that the deformations are relatively small and easier to track. Once motion between adjacent time frames has been estimated, it may be successively incorporated into a direct estimation of each frame to a single reference frame.

isolated time frames in the 3D implementation of this algorithm that the algorithm sometimes has problems converging to a well-matched deformation when the shape of the heart changes dramatically between two time frames. As one might expect, it is much easier to estimate the motion between two adjacent time frames where the motion is small than it is to estimate the motion between remote time frames, for example between end-diastole and end-systole. Perhaps this is partially due to the use of the implicit linear approximation in the Taylor expansion for the motion field. For large deformations, the Taylor expansion may no longer be a valid approximation. Another problem confounding the motion estimation is that the linear elastic strain energy model assumes small, infinitesimal motions. The calculated strain energy using the linear model is not as good a description of the material properties if large deformations are present [31].

The difficulty of obtaining a good estimation of the motion field for large deformations based on only two time frames is the principal impetus for investigating a 4D approach. As depicted in Fig. 1, it is known that motion of any voxel corresponding to cardiac tissue will be a cyclic, smooth motion, and the trajectory of the voxel described by the Lagrangian motion field should be a continuous, closed path. Because of this, we may first solve the easier problem of estimating the small motion between adjacent time frames, i.e., $\boldsymbol{m}(\boldsymbol{r}, t_i, t_{i+1})$, and then propagate these incremental motion fields to obtain the motion fields with respect to the reference frame, $\boldsymbol{m}(\boldsymbol{r}, t_1, t_i)$.

Implementation of the algorithm is therefore a nested

series of loops. Within the innermost loop, the conjugate gradient iteration is used to calculate an optimal $\delta \boldsymbol{m}(\boldsymbol{r}, t_i, t_j)$ between two time frames. The next loop level is over all time frames. Because a single calculation of $\delta \boldsymbol{m}(\boldsymbol{r}, t_i, t_j)$ does not usually obtain an adequate motion field, the calculation of the delta motion field over all times must be performed iteratively. Typically, 5-15 iterations are required for the conjugate gradient loop, and also 5-15 calculations of the outer loop are required before the algorithm converges. Finally, the algorithm alternates between computation of motion fields for adjacent time frames and computation of motion fields all pointing to the reference. The algorithm is initialized with a zero motion field between all time frames and the adjacent motion field is first computed. These incremental motion fields are used to initialize the composite motion fields all pointing to the reference frame. For example, in order to estimate the motion field between time frames 4 and 1, a composite initialization motion field for $\boldsymbol{m}(\boldsymbol{r}, t_1, t_4)$ is given by

$$\boldsymbol{m}(\boldsymbol{r}, t_1, t_4) = \boldsymbol{m}(\boldsymbol{r}, t_3, t_4) + \boldsymbol{m}(\boldsymbol{r}, t_2, t_3) + \boldsymbol{m}(\boldsymbol{r}, t_1, t_2). \quad (5)$$

Likewise, computation of the adjacent time motion fields from the distant ones is carried out in a similar manner. Alternating between the two forms of motion fields during the iteration insures consistency between the two. Note that when adding motion fields, the fields cannot simply be added as one would normally add two matrices of the same size. Rather, the motion fields must be forward projected to insure correct vector addition.

3. 3. Forward Sampling and Volume Summing

In an application where the deformed volumes are being summed together to form a composite volume, care must be taken during transformation sampling so that the composite volume represents a sum of all voxels from the source volumes. Many deformation algorithms make use of an Eulerian description of the motion field, which describes a particle's motion with respect to its final position in the deformation volume. This type of motion field representation requires a backward sampling scheme, typically where the value of each voxel in the deformed volume, $\hat{f}(\boldsymbol{r}) = f(\boldsymbol{r} - \boldsymbol{m})$ is obtained by a trilinear sampling at the non-discrete location, $\boldsymbol{r} - \boldsymbol{m}$. Such backward sampling does not guarantee that each voxel in the source volume will contribute to the deformation volume. We employ a forward sampling scheme [32] which distributes each voxel value of the source volume using a normalized Gaussian weighting in a single-pass calculation of the deformation as follows:

$$\hat{f}(\tilde{\boldsymbol{r}}) = \left(\sum_{\boldsymbol{r} + \boldsymbol{m} \in R(\tilde{\boldsymbol{r}})} \gamma_r f(\boldsymbol{r}) \right) \Big/ \left(\sum_{\boldsymbol{r} + \boldsymbol{m} \in R(\tilde{\boldsymbol{r}})} \gamma_r \right) \quad (6)$$

In this formulation, $R(\tilde{\boldsymbol{r}})$ represents the region where a deformed voxel would contribute via trilinear interpolation to a voxel at location, $\tilde{\boldsymbol{r}}$, and γ_r represents the interpolation weighting factor. Sampling in this manner insures that every voxel from all source volumes will contribute to the composite volume. Once each source volume has been deformed via the forward sampling technique, a motion corrected composite volume is formed by simple summing over all time frames.

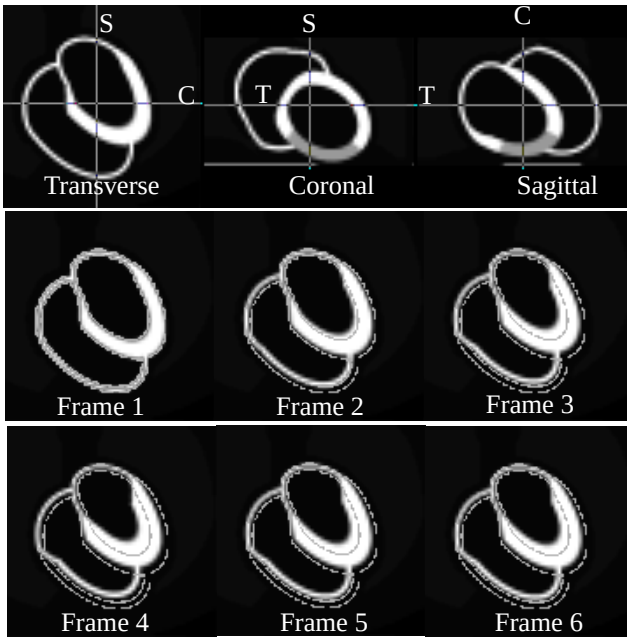

Fig. 2. Ellipsoidal Phantom with six time frames. Top row shows three orthogonal slices through the reference time frame 1. For reference, the intersection of each orthogonal plane with another is labeled as (T,C or S). Bottom two rows shows the transverse slice for all time frames. Edges of the reference frame are shown on all times to show the motion. The thick ellipsoid represents the left ventricle; the thin ellipsoids represent the right ventricle and atria.

4. RESULTS

The 4D algorithm was testing using two datasets. The first of these is a mathematical cardiac phantom based on ellipsoidal building blocks forming the human torso [33]. The second dataset is from PET data obtained during a gated cardiac study of a human subject. Fig. 2 shows the ellipsoidal phantom. A 60x60x37x6 voxel dataset was used to represent PET images of the heart at six time frames ranging from end-diastole to end-systole, and then back again.

Fig. 3 compares the results of deforming end-systole (frame 4) to match end-diastole (frame 1) using the 4D

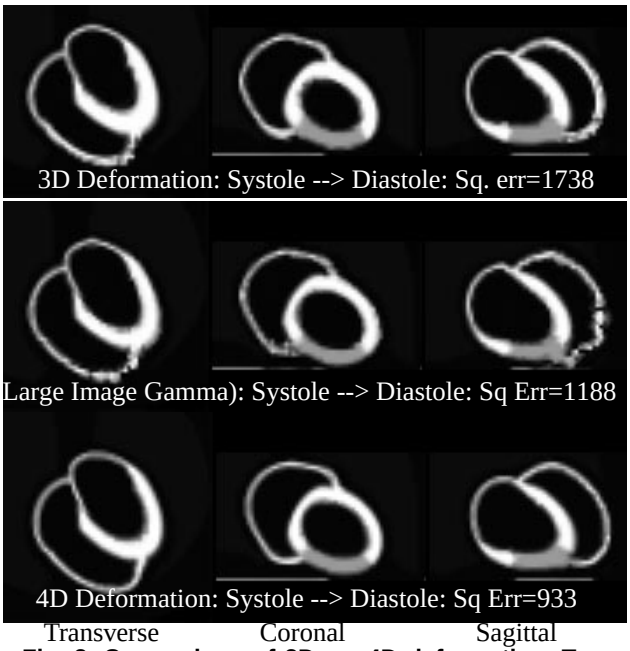

Fig. 3. Comparison of 3D vs. 4D deformation. Two attempts at deforming frame 4 to frame 1 using a 3D technique are only partially successful. The 4D technique using information from adjacent time frames was able to produce superior results.

technique verses a 3D technique. The top and bottom rows of images compare the 3D and 4D technique where the image matching weighting factor, γ_I, and the elastic material parameters were set to the same value for both simulations. Comparing with the three views of the reference volume shown in Fig. 2, it is seen that though both datasets deformed in roughly the correct direction (i.e. expansion), the 3D simulation had considerably more error. The deformed heart appears closer to systole and is more noisy for the 3D case. This is also indicated quantitatively by comparing the sum of squares difference between the reference and the deformation volumes for the 3D and 4D cases.

Another attempt using a considerably larger image matching weighting factor for the 3D case is shown in the middle row of Fig. 3. Here it is seen that even in the absence of noise, the algorithm cannot be forced to properly deform a source volume just by increasing the cost of image matching difference. Instabilities arise during the minimization process which can lead to artifacts, even though the overall sum of squares difference is less than for the small image gamma case.

Fig. 4 shows the desired result for the motion compensation algorithm in the noise free case. The first two rows of images show a simple sum of the six time frames uncorrected for motion blur, and a difference image

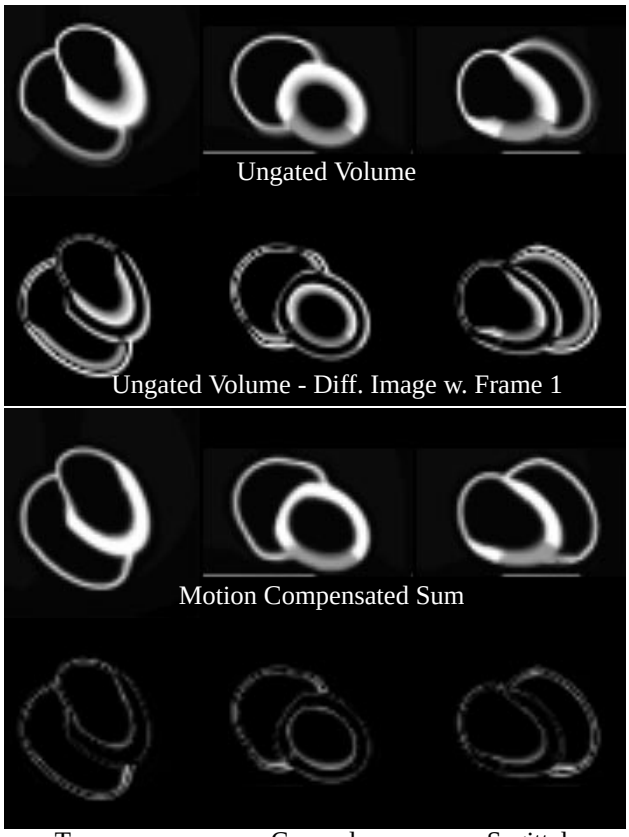

Fig. 4. Comparison of ungated phantom data to motion compensated sum. Top two rows show a direct sum of the six time frames and a difference image between the sum and the reference volume. The bottom two rows show the motion compensated sum and corresponding difference image. Blurring due to cardiac contractile motion is decreased in the motion compensated sum.

between the sum and the end diastole time frame. Because this is a simple sum of the original time frames, the wall thickness is blurred proportionally to the overall motion, effectively creating a blurred thickness that is the sum of the diastole and systole images. In contrast, the motion compensated sum and its corresponding difference image are seen in the bottom two rows. Though there are slight differences between the motion compensated volume and the reference gate, the algorithm was able to deform all the time frames so that they matched the reference quite closely.

The simulation using the mathematical phantom did not include noise, so there was no improvement in the image quality of the single end-diastole time frame when compared with the motion corrected sum. To show how the motion compensation technique can improve image quality, we can examine the algorithm's performance using real PET data from a human subject. Fig. 5 shows

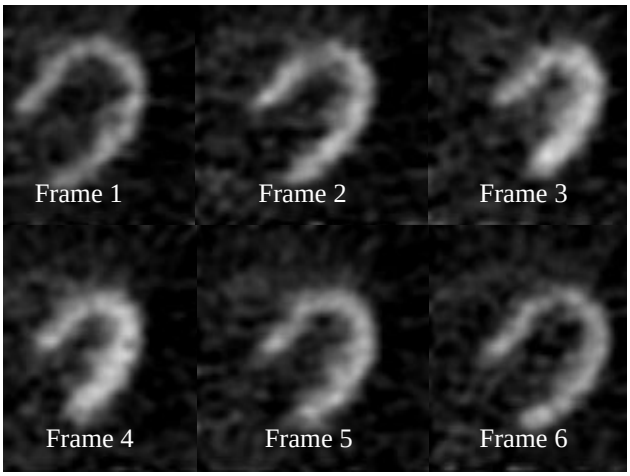

Fig. 5. First six 100 msec time frames from a human gated cardiac PET. The left ventricle is the principle feature in each transverse slice. End diastole is seen at frame 1. For motion compensation, all time frames must be deformed to match the shape of the heart at frame 1.

transverse slices through the left ventricle for the first six time frames of a cardiac dataset that was divided into nine 100 msec gates. In order to recognize the features in these data, the images have been smoothed considerably in the spatial domain. The transverse slice in the figure shows how the left ventricular walls contract and thicken as the heart cycle progresses from end-diastole in frame 1 to end-systole in frame 4. To produce a motion compensated composite volume, all volumes must be deformed so that the shape of the left ventricle appears as it does in frame 1.

Fig. 6 gives an example of the superior performance of the 4D algorithm over the 3D technique. The top row of the figure shows transverse, coronal and sagittal slices

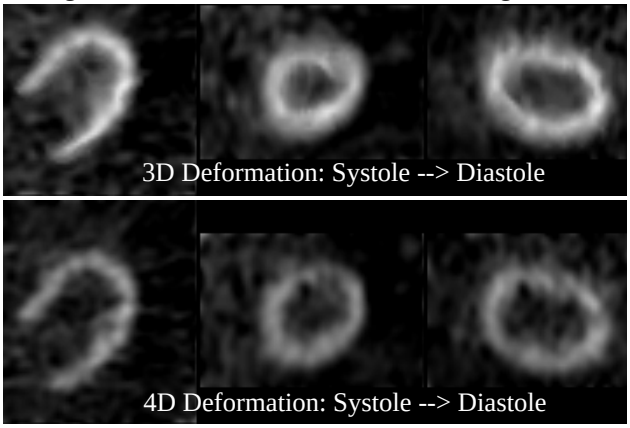

Fig. 6. Comparison of 3D vs. 4D deformations for real PET data. Attempts at deforming the systole volume to match the reference using only two time frames produced artifacts (top). In contrast, the technique using all time frames was able to produce a more faithful deformation (bottom).

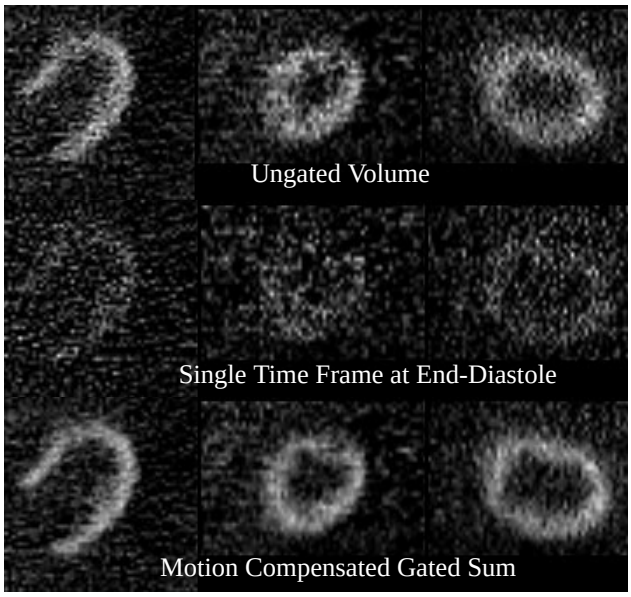

Fig. 7. Motion compensation results. Ungated data (top) suffers from blurring due to contractile motion. A single time frame at the same level of spatial smoothing has much less motion blur (middle), but is very noisy due to lack of statistics. The compensated sum (bottom) has less motion blur and better contrast to noise characteristics.

through a deformed volume where only time frames 1 and 3 were used in the calculation. Because of the large deformations necessary to match corresponding voxels in these two volumes, it appears that the algorithm had trouble estimating the true motion field, and thus produced a deformed volume with artifacts. In contrast, the bottom row of images shows the corresponding slices through a deformed volume that had been calculated using the full 4D information. Though we have no gold standard in this case, the deformed volume appears to be a much more reasonable estimate.

An example of the results of motion compensation for this human study are seen in Fig. 7. The images shown here have been smoothed considerably less in the spatial domain than the data seen in Figs. 5 and 6. The top row of images shows three orthogonal slices through an unprocessed sum of all nine gates, representing the way the data would look for an ungated PET acquisition. Here it is seen that the walls of the heart appear blurred due to the contractile motion. As one would expect, the resulting image is an average of the cardiac shape between all time frames between end-diastole and end-systole. The middle row of images shows a single time frame at end-diastole spatially smoothed to the same extent as the ungated sum. Though there is less motion blur in this single time frame, the lack of PET statistics results in a reconstructed image with much noise, with probably no better clinical value than the

ungated data. The motion compensated dataset is shown in the bottom row. Here, the 4D technique was used to estimate the motion between each time frame and the reference frame. All volumes were then deformed and summed. It is seen that the motion compensated volume has better image quality due to inclusion of more PET data as well as considerably less motion blur.

5. Discussion and Summary

Results presented here indicate that the characteristic of temporal continuity in a 4D cardiac PET dataset may be exploited to improve the performance of a motion estimation algorithm. In our experience, the principal advantage to the 4D technique when compared to a technique using only two time frames is that the ability to more accurately estimate large deformations is dramatically improved. We have found that for the 3D case, estimation of motion between greatly differing time frames is possible, for example, a direct estimation of the motion between end diastole and end systole; however, selection of the image matching and elastic material parameters is much more sensitive.

Temporal continuity is enforced by two methods in this algorithm. First is by means of a prediction motion field, which makes the rather simple assumption of constant particle velocity between time frames. Second, temporal continuity is exploited by the use of concatenated incremental motion fields that initialize the estimation of motion fields between distant time frames. The prediction fields augment the algorithm's ability to estimate motion between noisy adjacent time frames, and the concatenated incremental motion fields allow improved estimation for large deformations.

In its current form, the prediction field calculation is a crude approximation, since it assumes constant velocity between time frames. For a cyclic motion like the motion of the heart, the deformation is not described by a constant velocity, but instead actually reverses in direction. Indeed, we have found that for the noise-free case, the prediction field can actually degrade performance in time frames where the particle velocity is changing, for example, in the time frames before and after end-systole. However, for noisy data, we have found that even this rough estimate of a prediction field is beneficial. Future work will focus on more accurate prediction models, such as the Kalman filtering or Fourier techniques that have been used successfully in a similar motion tracking application [19,20, 21].

It is acknowledged that the results presented here are almost exclusive qualitative in nature. As is often the case when using real PET data obtained from a human subject, comparison with a gold standard is difficult. The mathematical phantom allows for a more controlled comparison between truth and the algorithm's results; however, the ellipsoidal description of this phantom makes it difficult to assess the exact motion vectors between corresponding points. Work is underway to develop a torso phantom with a parametric description that would allow more detailed comparison. We have used such parametrically described phantoms in the past [32], but these have been only isolated left ventricle models and perhaps are not as good a test as a full torso model. It is also important to test the algorithm on a phantom in a noise environment consistent with PET reconstructions.

Based on arguments regarding the aperture problem in optical flow, and recognizing the limited spatial and temporal resolution of even the best cardiac PET scanners, it may be debated whether any algorithm could be expected to follow the complex cardiac motion for small features within the cardiac walls. An overall goal of this algorithm is to provide more spatially and temporally resolved images of the heart, so that clinicians may better identify cardiac defects that would be seen as nonuniformities in radiotracer uptake. If there are no defects, a PET image of the left ventricle is relatively featureless. On the other hand, if a small defect does exist, it may be too noisy in a single time frame or too blurred in a spatially or temporally smoothed volume for a motion estimation algorithm to track. However, results using the ellipsoidal phantom indicate that when combined with an adequate material model regularization, the feature space of even this smooth phantom is rich enough for adequate tracking. Therefore, the "aperture problem" as noted in the 2D optical flow community may not be problematic here. It is hoped that future testing on more realistic phantoms will answer whether the algorithm can adequately track small features. At this point though, we can at least say that the motion compensation algorithm *reduces* motion blur. It is hoped that future algorithm improvements will provide even more accurate motion estimations.

6. Acknowledgments

This work was supported by the National Heart, Lung and Blood Institute of the U.S. Dept. of Health and Human Services under grant P01-HL25840; and by the U.S. Dept. of Energy under Contract DE-AC03-76SF00098.

7. References

[1] S. M. Song and R. M. Leahy. Computation of 3-D velocity fields from 3-D cine CT images of a human heart. *IEEE Trans Med Imag*, 10(1):295–306, 1991.

[2] Z. Zhou, C. E. Synolakis, R. M. Leahy, and S. M. Song. Calculation of 3D internal displacement fields from 3D X-ray computer tomographic images. *Proc R Soc Lond A*,

449(1937):537–554, 1995.

[3] B. Horn and B. G. Schunck. Determining optical flow. *Artificial Intell.*, 17:185–203, 1981.

[4] R. Bajcsy and S. Kovacic. Multiresolution elastic matching. *Comput. Vision, Graph., Image Proc.*, 46:1–21, 1989.

[5] J. Gee, M. Reivich, and R. Bajcsy. Elastically deforming a 3D atlas to match anatomical brain images. *J Comput Assist Tomogr*, 17(2):225–236, 1993.

[6] G. E. Christensen, R. D. Rabbitt, and M. I. Miller. 3d brain mapping using a deformable neuroanatomy. *Phys Med Biol*, 39:609–618, 1994.

[7] C. Davatzikos. Spatial normalization of 3D brain images using deformable models. *J Comput Assist Tomogr*, 20(4):656–665, 1996.

[8] M. H. Davis, A. Khotanzad, D. P. Flamig, and S. E. Harms. A physics-based coordinate transformation for 3-D image matching. *IEEE Trans Med Imag*, 16(3):317–328, 1997.

[9] P. Thompson and A. W. Toga. A surface-based technique for warping three-dimensional images of the brain. *IEEE Trans Med Imag*, 15(4):402–417, 1996.

[10] J. Talairach and P. Tournoux. *Co-planar Stereotactic Atlas of the Human Brain: 3-Dimensional Proportional System: An Approach to Cerebral Imaging*. Georg Theime Verlag, Stuttgart, 1988.

[11] R. P. Woods, S. T. Grafton, J. D. Watson, N. L. Sicotte, and J. C. Mazziotta. Automated image registration: II. intersubject validation of linear and nonlinear models. *J Comput Assist Tomogr*, 22(1):153–165, 1998.

[12] F. L. Bookstein. Principal warps: Thin-plate splines and the decomposition of deformations. *IEEE Trans Patt Anal Machine Intell*, 11(6):567–585, 1989.

[13] J. Declerck, G. Subsol, J. Thirion, and N. Ayache. Automatic retrieval of anatomical structures in 3D medical images. In *Computer Vision, Virtual Reality and Robotics in Medicine*, vol 905, pp 153–162. Springer-Verlag, INRIA, Sophia-Antipolis, France, 1995.

[14] D. Heeger. Optical flow using spatiotemporal filters. *Int J Comp Vision*, 1:279–302, 1988.

[15] M. J. Black and P. Anandan. Robust dynamic motion estimation over time. In *Proc IEEE CVPR*, pages 296–302, Maui, Hawaii, 1991.

[16] D. J. Fleet and K. Langley. Recursive filters for optical flow. *IEEE Trans Patt Anal Machine Intell*, 17(1):61–67, 1995.

[17] P. Shi, A. Amini, G. Robinson, et. al. Shape-based 4d left ventricular myocardial function analysis. In *IEEE WBIA*, pp 88–97, Seattle, WA, 1994.

[18] J. Declerck, J. Feldmar, and N. Ayache. Definition of a 4D continuous polar transformation for the tracking and the analysis of LV motion. Report 2105, INRIA Sophia-Antipolis France, 1996.

[19] F. G. Meyer, T. Constable, A. J. Sinusas, and J. S. Duncan. Tracking myocardial deformation using phase contrast MR velocity fields: A stochastic approach. *IEEE Trans Med Imag*, 4(453-465), 15.

[20] Y. Zhu and N. J. Pelc. A spatiotemporal model of cycle kinematics and its application to analyzing nonrigid motion with MR velocity images. *IEEE Trans Med Imag*, 18(7):557–569, 1999.

[21] Y. Zhu and N. J. Pelc. Three-dimensional motion tracking with volumetric phase contrast MR velocity imaging. *J Magn Reson Imag*, 9:111–118, 1999.

[22] J. Huang, D. Abendschein, V. Davila-Roman, and A. A. Amini. Spatio-temporal tracking of myocardial deformations with a 4d b-spline model from tagged MRI. *IEEE Trans Med Imag*, 18(10):957–972, 1999.

[23] T. L. Faber, E. M. Stokely, et. al. A model-based four-dimensional left ventricular surface detector. *IEEE Trans Med Imag*, 10(3):321–329, 1991.

[24] T. L. Faber, C. D. Cooke, et. al.. Left ventricular function and perfusion from gated SPECT perfusion images: An integrated method. *J Nucl Med*, 40(4):650–659, 1999.

[25] G. Germano, H. Kiat, et. al. Automatic quantification of ejection fraction from gated myocardial perfusion SPECT. *J Nucl Med*, 36:2138–2147, 1995.

[26] G. J. Klein. Forward deformation of PET volumes using non-uniform elastic material constraints. In A. Kuba and J. Attila, ed, *IPMI*, vol 1613, pp 358–363. Springer, 1999.

[27] G. J. Klein. Deformable models for volume feature tracking. Report LBNL-43257, Lawrence Berkeley National Laboratory, 1999. (Ph.D. Thesis).

[28] P. Nielsen, I. Le Grice, B. H. Smaill, and P. J. Hunter. Mathematical model of geometry and fibrous structure of the heart. *Am. J. Physiol.*, 29:H1365–H1378, 1991.

[29] P. J. Hunter, M. P. Nash, and G. B. Sands. Computational electromechanics of the heart. In A. V. Panfilov and A. V. Holden, editors, *Computational Biology of the Heart*, ch 12, pp 345–404. John Wiley and Sons, 1997.

[30] B. K. P. Horn. *Robot Vision*. The MIT Press, Cambridge, Massachusetts, 1986.

[31] L. K. Waldman, Y. C. Fung, and J. W. Covell. Transmural myocardial deformation in the canine left ventricle. *Circulation Research*, 57(1):152–163, 1985.

[32] G. J. Klein. Forward deformation of PET volumes using material constraints. In B. Vemuri, ed, *WBIA, Santa Barbara, CA*, pp 64–71. IEEE Computer Society, 1998.

[33] P. H. Pretorius, W. Xia, et. al. Evaluation of right and left ventricular volume and ejection fraction using a mathematical cardiac torso phantom. *J Nucl Med*, 38(10):1528–1535, 1997.